

## EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	8549	707/200-204.ccls.	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 13:43
L4	9	I1 and (identifier and first and second and third and state and replication).clm.	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 13:49
L5	2	I1 and (identifier and first and second and third and state and replication).clm. and @ad<"20021127"	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 13:46
L7	0	I1 and (identifier and first and second and third and state and replication and lock\$3).clm. and @ad<"20021127"	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 13:49
L8	854	711/152.ccls.	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 13:49
L9	0	I8 and (identifier and first and second and third and state and replication and lock\$3).clm. and @ad<"20021127"	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 13:49

## EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S20 2	423	"data mover" same storage	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S20 3	116	"data mover" near5 storage and @ad<"20021127"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S20 4	96	(backup replicat\$3) near2 (flag status indicator) and data and state and identifier and lock	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S20 5	109	("first state" same "second state" same "third state") and data and identifier and lock\$3 and @ad<"20021127"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S20 6	8	"first state" and "second state" and "third state" and data and identifier and (logical near2 lock)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S20 7	38	("4682305"   "4985863"   "5161214"   "5187750"   "5287497").PN. OR ("5408630").URPN.	US-PGPUB; USPAT; USOCR	OR	ON	2008/01/02 11:51
S20 8	6	((("20040143745") or ("6850938") or ("6857068")).PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/01/02 11:51
S20 9	70	("5001628"   "5032979"   "5113519"   "5325524"   "5367698"   "5408619"   "5408657"   "5418966"   "5434974"   "5434994"   "5459860"   "5463774"   "5491817"   "5519855"   "5608903"   "5613079"   "5677851"   "5689706").PN. OR ("5832487").URPN.	US-PGPUB; USPAT; USOCR	OR	ON	2008/01/02 11:51

## EAST Search History

S21 0	181	atom\$3 near2 transaction with process\$3	US-PGPUB; USPAT; USOCR	OR	ON	2008/01/02 11:51
S21 1	1801	atom\$3 with transaction	US-PGPUB; USPAT; USOCR	OR	ON	2008/01/02 11:51
S21 2	1	("7225302").PN.	US-PGPUB; USPAT; EPO	OR	OFF	2008/01/02 12:06
S21 3	6	((("5522077") or ("7225302") or ("7143414"))).PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/01/02 11:51
S21 4	72	atom\$3 near2 transaction with process\$3 and (identifier ID) and @ad<"20021127"	US-PGPUB; USPAT; USOCR	OR	ON	2008/01/02 11:51
S21 5	64230	"707"/\$.ccls.	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 11:51
S21 6	207	((backup replicat\$3) near2 commit\$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S21 7	167	(storage near4 controller) and (metadata same (lock locking))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S21 8	521	multiple with semaphore	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 11:51
S21 9	120	atom\$3 near2 transaction with process\$3 and @ad<"20021127"	US-PGPUB; USPAT; USOCR	OR	ON	2008/01/02 11:51
S22 0	237	((backup replicat\$3) near4 (flag status)) and commit\$3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S22 1	118	atom\$3 near2 transaction with process\$3 and @ad<"20021117"	US-PGPUB; USPAT; USOCR	OR	ON	2008/01/02 11:51

## EAST Search History

S22 2	216	"data mover" near5 storage	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S22 3	49	atom\$3 near2 transaction with process\$3 and identifier and @ad<"20021127"	US-PGPUB; USPAT; USOCR	OR	ON	2008/01/02 11:51
S22 4	113	((multiple near2 state) same lock)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S22 5	32	((backup replicat\$3) near4 (flag status)) same commit\$3) and (ID identifier GUID)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S22 6	11	(storage near4 controller) same metadata same lock	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S22 7	41	((multiple near2 state) near4 lock)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S22 8	37	lock near2 share near2 exclusive	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S22 9	61	(lock\$3 near2 id) near4 state	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S23 0	8	((("20040153480") or ("6029178") or ("6032158") or ("20040148447"))). PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2008/01/02 11:51

## EAST Search History

S23 1	2	lock\$3 same (metadata) same group same (backup replicat\$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S23 2	10	tagged near2 backup	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S23 3	72	(atom\$3 near2 transaction) with process\$3 and (identifier ID) and @ad<"20021127"	US-PGPUB; USPAT; USOCR	OR	ON	2008/01/02 11:51
S23 4	35	(multiple with semaphore) same (copy replicat\$3)	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 11:51
S23 5	6	(multiple and semaphore).ti. and (semaphore lock status)	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 11:51
S23 6	18	(multiple with semaphore) same (copy replicat\$3)and @ad<"20021127"	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 11:51
S23 7	2	(atom\$3 near2 transaction) with process\$3 same (identifier ID) and @ad<"20021127"	US-PGPUB; USPAT; USOCR	OR	ON	2008/01/02 11:51
S23 8	30	(time near stamp) same ("globally unique identifier" GUID) and database and @ad<"20021127"	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 11:51
S23 9	10	((time near stamp) and ("globally unique identifier" GUID)) and (object near2 oriented) and "shared memory" and @ad<"20021127"	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 11:51
S24 0	4	((time near stamp) same ("globally unique identifier" GUID)) and (object near2 oriented) and memory and @ad<"20021127"	US-PGPUB; USPAT; EPO	OR	ON	2008/01/02 11:51
S24 1	160	((backup replicat\$3 duplicat\$3) near2 commit\$3) and transaction	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51

## EAST Search History

S24 2	173	((backup replicat\$3) near4 (flag status)) and commit\$3 and (ID identifier GUID)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51
S24 3	129	(storage near2 controller) and (metadata same (lock locking))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/02 11:51

[Web](#) [Images](#) [Maps](#) [News](#) [Shopping](#) [Gmail](#) [more ▾](#)[Sign in](#)

Google

[Advanced Search](#)  
[Preferences](#)

Web

Results 1 - 10 of about 23,000 for **multiple state "data object" replication lock**. (0.28 seconds)**Method and software application for avoiding data loss - US Patent ...**

A method of avoiding data loss in a **data object replication** process, ... yet ready for **replication**; setting a shared **lock** on the electronic data element; ...

[www.patentstorm.us/patents/7225302-claims.html](http://www.patentstorm.us/patents/7225302-claims.html) - 24k - [Cached](#) - [Similar pages](#)

**Enterprise Software Pattern Synopses**

**Read/Write Lock**; The Read/Write **lock** pattern is a specialized form of the Scheduler pattern. It allows **multiple** clients to concurrently have read access to ...

[www.mindspring.com/~mgrand/pattern\\_synopses3.htm](http://www.mindspring.com/~mgrand/pattern_synopses3.htm) - 55k - [Cached](#) - [Similar pages](#)

**OCRed document**

of **replication** because a remote **lock** request may be neces- ... **multiple** readers may exist if an object is in the **state**. shared, we define our.

[ieeexplore.ieee.org/iel5/3016/8569/00375489.pdf?arnumber=375489](http://ieeexplore.ieee.org/iel5/3016/8569/00375489.pdf?arnumber=375489) - [Similar pages](#)

**Parallel Computing : On the design of global object space for ...**

The JMM defines memory consistency semantics of **multi**-threaded Java programs. There is a **lock** associated with each object in Java. Based on the JMM proposed ...

[linkinghub.elsevier.com/retrieve/pii/S0167819103001467](http://linkinghub.elsevier.com/retrieve/pii/S0167819103001467) - [Similar pages](#)

**Enterprise Software Pattern Synopses**

**Object Replication**; The Object **Replication** pattern addresses issues related to keeping the **state** of the client and server objects consistent. ...

[www.markgrand.com/pattern\\_synopses3.htm](http://www.markgrand.com/pattern_synopses3.htm) - 54k - [Cached](#) - [Similar pages](#)

**Concurrency versus availability: atomicity mechanisms for ...**

The author defines a replicated object as a typed **data object** whose **state** is stored at **multiple** locations to enhance availability. A **replication** method is ...

[portal.acm.org/citation.cfm?id=27643](http://portal.acm.org/citation.cfm?id=27643) - [Similar pages](#)

**Overview**

optional data **replication** between **multiple** SOSS stores running in **multiple** ... and update any stored **data object** from any participating server in the store. ...

[www.scaleoutsoftware.com/support/stateServer/sooss\\_help/introduction/overview.htm](http://www.scaleoutsoftware.com/support/stateServer/sooss_help/introduction/overview.htm) - 42k -

[Cached](#) - [Similar pages](#)

**LNCS 2848 - Efficient Replication of Large Data Objects**

Let  $x$  be the **data object** we are replicating.  $x$  takes values in a set  $V$ , and has default value  $v$ .  $0 \leq x$  can be read and written to. To replicate **multiple** ...

[www.springerlink.com/index/3A3FG4DHEJD78PXX.pdf](http://www.springerlink.com/index/3A3FG4DHEJD78PXX.pdf) - [Similar pages](#)

**Patents in Class 707/8**

A **multi**-protocol **lock** manager efficiently manages granting, revoking and releasing of various ... A method or apparatus for cooperative data **replication**. ...

[www.freepatentsonline.com/CCL-707-8.html](http://www.freepatentsonline.com/CCL-707-8.html) - 77k - [Cached](#) - [Similar pages](#)

**File Or Database Maintenance - Coherency (e.g., Same View To ...**

08/24/06 - 20060190497 - Support for schema evolution in a **multi**-node peer-to-peer **replication** environment A method, system and article of manufacture of ...

[www.freshpatents.com/x1707201000psbc.php](http://www.freshpatents.com/x1707201000psbc.php) - 65k - [Cached](#) - [Similar pages](#)



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide


**THE ACM DIGITAL LIBRARY**

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used: **replication multiple state data object**Found **138,924** of **216,412**

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)

Display results


[Search Tips](#)
[Try this search in The ACM Guide](#)
☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

### 1 [The effect of failure and repair distributions on consistency protocols for replicated data objects](#)

John L. Carroll, Darrell D. E. Long

March 1989 **Proceedings of the 22nd annual symposium on Simulation ANSS '89**

Publisher: IEEE Computer Society Press

Full text available: pdf(1.26 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The accessibility of vital information can be enhanced by replicating the data on several sites, and employing a consistency control protocol to manage the copies. Various protocols have been proposed to ensure that only current copies of the data can be accessed. The effect these protocols have on the accessibility of the replicated data is investigated by simulating the operation of the network and measuring the performance. Several strategies for replica maintenance are consid ...

### 2 [Concurrency versus availability: atomicity mechanisms for replicated data](#)



Maurice Herlihy

August 1987 **ACM Transactions on Computer Systems (TOCS)**, Volume 5 Issue 3

Publisher: ACM Press

Full text available: pdf(2.02 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

A replicated object is a typed data object that is stored redundantly at multiple locations to enhance availability. Most techniques for managing replicated data have a two-level structure: At the higher level, a replica-control protocol reconstructs the object's state from its distributed components, and at the lower level, a standard concurrency-control protocol synchronizes accesses to the individual components. This paper explores an alternative approach to managing replicated data by p ...

### 3 [DistView: support for building efficient collaborative applications using replicated objects](#)



Atul Prakash, Hyong Sop Shim

October 1994 **Proceedings of the 1994 ACM conference on Computer supported cooperative work CSCW '94**

Publisher: ACM Press

Full text available: pdf(1.61 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The ability to share synchronized views of interactions with an application is critical to supporting synchronous collaboration. This paper suggests a simple synchronous collaboration paradigm in which the sharing of the views of user/application interactions occurs at the window level within a multi-user, multi-window application. The paradigm is



incorporated in a toolkit, DistView, that allows some of the application windows to be shared at a fine-level of granularity, while still allowin ...

**Keywords:** active objects, collaboration technology, concurrency control, distributed objects, groupware, multiuser interfaces, replicated objects, shared windows

#### 4 A quorum-consensus replication method for abstract data types



Maurice Herlihy

February 1986 **ACM Transactions on Computer Systems (TOCS)**, Volume 4 Issue 1

**Publisher:** ACM Press

Full text available: [pdf\(1.66 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Replication can enhance the availability of data in distributed systems. This paper introduces a new method for managing replicated data. Unlike many methods that support replication only for uninterpreted files, this method systematically exploits type-specific properties of objects such as sets, queues, or directories to provide more effective replication. Each operation requires the cooperation of a certain number of sites for its successful completion. A quorum for an operation is any s ...

#### 5 Concurrency and availability as dual properties of replicated atomic data



M. Herlihy

April 1990 **Journal of the ACM (JACM)**, Volume 37 Issue 2

**Publisher:** ACM Press

Full text available: [pdf\(1.77 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

A replicated data object is a typed object that is stored redundantly at multiple locations in a distributed system. Each of the object's operations has a set of quorums, which are sets of sites whose cooperation is needed to execute that operation. A quorum assignment associates each operation with its set of quorums. An operation's quorums determine its availability, and the constraints governing an object's quorum assignments determine the range of availability properties realizable by r ...

#### 6 Mobility: Flexible on-device service object replication with replets



Dong Zhou, Nayeem Islam, Ali Ismael

May 2004 **Proceedings of the 13th international conference on World Wide Web WWW '04**

**Publisher:** ACM Press

Full text available: [pdf\(887.11 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

An increasingly large amount of Web applications employ service objects such as Servlets to generate dynamic and personalized content. Existing caching infrastructures are not well suited for caching such content in mobile environments because of disconnection and weak connection. One possible approach to this problem is to replicate Web-related application logic to client devices. The challenges to this approach are to deal with client devices that exhibit huge divergence in resource ...

**Keywords:** capability, preference, reconfiguration, replication, service, synchronization

#### 7 A meta model and an infrastructure for the non-transparent replication of object databases



Werner Dreyer, Klaus R. Dittrich

November 2000 **Proceedings of the ninth international conference on Information and knowledge management CIKM '00**

**Publisher:** ACM Press

Full text available: [pdf\(179.36 KB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

**Keywords:** object databases, object replication, replication meta models

8 A task- and data-parallel programming language based on shared objects



Saniya Ben Hassen, Henri E. Bal, Criel J. H. Jacobs

November 1998 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 20 Issue 6

**Publisher:** ACM Press

Full text available: pdf(434.44 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Many programming languages support either task parallelism, but few languages provide a uniform framework for writing applications that need both types of parallelism or data parallelism. We present a programming language and system that integrates task and data parallelism using shared objects. Shared objects may be stored on one processor or may be replicated. Objects may also be partitioned and distributed on several processors. Task parallelism is achieved by forking processes remotely a ...

**Keywords:** data parallelism, shared objects, task parallelism

9 Research session: spatial and temporal databases: A trajectory splitting model for efficient spatio-temporal indexing

Slobodan Rasetic, Jörg Sander, James Elding, Mario A. Nascimento

August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05**

**Publisher:** VLDB Endowment

Full text available: pdf(400.37 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper addresses the problem of splitting trajectories optimally for the purpose of efficiently supporting spatio-temporal range queries using index structures (e.g., R-trees) that use minimum bounding hyper-rectangles as trajectory approximations. We derive a formal cost model for estimating the number of I/Os required to evaluate a spatio-temporal range query with respect to a given query size and an arbitrary split of a trajectory. Based on the proposed model, we introduce a dynamic progr ...

10 Consistency and replication: Application specific data replication for edge services



Lei Gao, Mike Dahlin, Amol Nayate, Jiandan Zheng, Arun Iyengar

May 2003 **Proceedings of the 12th international conference on World Wide Web WWW '03**

**Publisher:** ACM Press

Full text available: pdf(476.22 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The emerging edge services architecture promises to improve the availability and performance of web services by replicating servers at geographically distributed sites. A key challenge in such systems is data replication and consistency so that edge server code can manipulate shared data without incurring the availability and performance penalties that would be incurred by accessing a traditional centralized database. This paper explores using a distributed object architecture to build an edge s ...

**Keywords:** availability, data replication, distributed objects, edge services, performance, wide area networks (WAN)

11 Replicated data management in distributed database systems

Sang Hyuk Son

November 1988 **ACM SIGMOD Record**, Volume 17 Issue 4



**Publisher:** ACM Press

Full text available: pdf(835.25 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Replication is the key factor in improving the availability of data in distributed systems. Replicated data is stored at multiple sites so that it can be accessed by the user even when some of the copies are not available due to site failures. A major restriction to using replication is that replicated copies must behave like a single copy, i.e., mutual consistency as well as internal consistency must be preserved. Synchronization techniques for replicated data in distributed database systems ...

## 12 Optimistic replication



Yasushi Saito, Marc Shapiro

March 2005 **ACM Computing Surveys (CSUR)**, Volume 37 Issue 1

**Publisher:** ACM Press

Full text available: pdf(656.72 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Data replication is a key technology in distributed systems that enables higher availability and performance. This article surveys optimistic replication algorithms. They allow replica contents to diverge in the short term to support concurrent work practices and tolerate failures in low-quality communication links. The importance of such techniques is increasing as collaboration through wide-area and mobile networks becomes popular. Optimistic replication employs algorithms not seen in tradition ...

**Keywords:** Replication, disconnected operation, distributed systems, large scale systems, optimistic techniques

## 13 The state of the art in distributed query processing



Donald Kossmann

December 2000 **ACM Computing Surveys (CSUR)**, Volume 32 Issue 4

**Publisher:** ACM Press

Full text available: pdf(455.39 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Distributed data processing is becoming a reality. Businesses want to do it for many reasons, and they often must do it in order to stay competitive. While much of the infrastructure for distributed data processing is already there (e.g., modern network technology), a number of issues make distributed data processing still a complex undertaking: (1) distributed systems can become very large, involving thousands of heterogeneous sites including PCs and mainframe server machines; (2) the state ...

**Keywords:** caching, client-server databases, database application systems, dissemination-based information systems, economic models for query processing, middleware, multitier architectures, query execution, query optimization, replication, wrappers

## 14 Mobile computing and applications: A mechanism for replicated data consistency in mobile computing environments



José Maria Monteiro, Ângelo Brayner, Sérgio Lifschitz

March 2007 **Proceedings of the 2007 ACM symposium on Applied computing SAC '07**

**Publisher:** ACM Press

Full text available: pdf(148.47 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Mobile computing allows for the development of new and sophisticated database applications. Such applications require the reading of current and consistent data. In order to improve data availability, increase performance and maximize throughput, data replication is used. However, due to inherent limitations in mobile and other loosely-

coupled environments, the concurrency control and replica control mechanisms must be revisited. This paper proposes a new protocol that guarantees the consistency ...

**Keywords:** concurrency control, data replication, mobile computing

15 Apologizing versus asking permission: optimistic concurrency control for abstract data types



M. Herlihy

March 1990 **ACM Transactions on Database Systems (TODS)**, Volume 15 Issue 1

**Publisher:** ACM Press

Full text available: pdf(2.38 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

An optimistic concurrency control technique is one that allows transactions to execute without synchronization, relying on commit-time validation to ensure serializability. Several new optimistic concurrency control techniques for objects in decentralized distributed systems are described here, their correctness and optimality properties are proved, and the circumstances under which each is likely to be useful are characterized. Unlike many methods that classify operations only a ...

16 What service replication middleware can learn from object replication middleware



Johannes Osrael, Lorenz Frohofer, Karl M. Goeschka

November 2006 **Proceedings of the 1st workshop on Middleware for Service Oriented Computing (MW4SOC 2006) MW4SOC '06**

**Publisher:** ACM Press

Full text available: pdf(340.42 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Replication is a well-known technique to enhance dependability and performance in distributed systems. A plethora of replication middleware for distributed object systems has been proposed in the past decade. However, replication in service-oriented systems is still in its infancy. In this paper, we analyze some of the proposed service replication middleware solutions and compare them on an architectural level with object replication middleware. In particular, we focus on replication middleware ...

**Keywords:** architecture, distributed object systems, middleware, replication, service-oriented systems

17 Adaptive distributed data management with weak consistent replicated data



Richard Lenz

February 1996 **Proceedings of the 1996 ACM symposium on Applied Computing SAC '96**

**Publisher:** ACM Press

Full text available: pdf(1.07 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

**Keywords:** coherency control, consistency island, distributed data management, need-to-know principle, weak consistency replication

18 Replicated state space approach for parallel simulation



John B. Gilmer, Jung Pyo Hong

December 1986 **Proceedings of the 18th conference on Winter simulation WSC '86**

**Publisher:** ACM Press

Full text available: pdf(519.02 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Parallel processing offers the possibility of greatly increased performance for simulations

which are computationally bound on existing machines. On shared memory machines, such as the BBN Butterfly, a natural approach is to allocate entities to be processed on different processors with locks used to prevent synchronization problems for a state space in global memory. Parallel processors having local memory only, such as the hypercube architectures, cannot use this approach. Such machines a ...

### 19 Replicated objects in time warp simulations

 Divyakant Agrawal, Jonathan R. Agre  
December 1992 **Proceedings of the 24th conference on Winter simulation WSC '92**


**Publisher:** ACM Press

Full text available:  pdf(845.76 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

### 20 Mobile objects in distributed Oz

 Peter Van Roy, Seif Haridi, Per Brand, Gert Smolka, Michael Mehl, Ralf Scheidhauer  
September 1997 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 19 Issue 5

**Publisher:** ACM Press

Full text available:  pdf(484.83 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Some of the most difficult questions to answer when designing a distributed application are related to mobility: what information to transfer between sites and when and how to transfer it. Network-transparent distribution, the property that a program's behavior is independent of how it is partitioned among sites, does not directly address these questions. Therefore we propose to extend all language entities with a network behavior that enables efficient distributed program ...

**Keywords:** latency tolerance, mobile objects, network transparency

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2008 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

**Search Results**

[BROWSE](#)

[SEARCH](#)


[IEEE XPLORE GUIDE](#)

[SUPPORT](#)

Results for "((multiple state data object lock replication)<in>metadata )"

Your search matched 0 of 1715275 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

 [e-mail](#)  [printer](#)



» Search Options

[View Session History](#)

[New Search](#)

» Key

- IEEE JNL IEEE Journal or Magazine
- IET JNL IET Journal or Magazine
- IEEE CNF IEEE Conference Proceeding
- IET CNF IET Conference Proceeding
- IEEE STD IEEE Standard

Modify Search

((multiple state data object lock replication)<in>metadata )

[Search](#)

☐ Check to search only within this results set

Display Format: ☒ Citation ☐ Citation & Abstract

[IEEE/IET](#)

[Books](#)

[Educational Courses](#)

[Application Notes](#)

IEEE/IET journals, transactions, letters, magazines, conference proceedings, and standards.

[view selected items](#)

[Select All](#) [Deselect All](#)

**No results were found.**

Please edit your search criteria and try again. Refer to the Help pages if you need assistance revising your search.

[Help](#) [Contact Us](#) [Privacy & Security](#)

© Copyright 2007 IEEE – All Rights Reserved